arXiv:1404.4917v1 [math.PR] 19 Apr 2014

On the explanatory power of principal components

Daniel A. Díaz–Pachón[1]*, J. Sunil Rao[1]†and Jean-Eudes Dazard[2]‡
[1] University of Miami, Miami, FL, USA
[2] Case Western Reserve University, Cleveland, OH, USA

April 22, 2014

Abstract

We show that if we have an orthogonal base $(u_1, \ldots, u_p)$ in a p-dimensional vector space, and select $p+1$ vectors $v_1, \ldots, v_p$ and w such that the vectors traverse the origin, then the probability of w being to closer to all the vectors in the base than to $v_1, \ldots, v_p$ is at least $1/2$ and converges as p increases to infinity to a normal distribution on the interval [-1,1]; i.e., $\Phi(1) - \Phi(-1) \approx 0.6826$. This result has relevant consequences for Principal Components Analysis in the context of regression and other learning settings, if we take the orthogonal base as the direction of the principal components.

1 Introduction

Principal Components Analysis (PCA) has been popular since its inception due to its two main features: 1. Explaining well the variation of the information and 2. Because of the first one, being able to reduce the space to those variables that explain most of the variation. In the context of regression, PCA has been used to explain the response in terms of the Principal Components (PC) of the original input variables (see for instance, [8], p. 75; [6]; [9], p. 307; [10]). However, because one of the main features of PCA is dimension reduction (feature 2 mentioned above) it is not always clear that PCA can be used and to still get that the response is well explained by the reduced space of the main PC (see for instance [3], p. 72; [4]; [5]). In fact, as early as 1957, in one of the first references to PCA in regression, Hotelling was doubtful of the method because of this same reason ([7]). For a historical review of PCA in regression settings see [2].

*Ddiaz3@med.miami.edu
†jrao@biostat.med.miami.edu
‡jxd101@case.edu

Note the interesting fact that the earliest of our references advocating the use of PCA in regression ([8]) was proposed in the same year that its first criticism was known ([7]). The questionings of PCA in regression seem to be as old as the method itself. Now, the method and its questionings transcend regression going further to more general learning settings. But whether regression or any other learning setting, the two interesting questions to ask in order to determine the usefulness of the method are the following: 1. Is it the case that (most of the times) the response variable is better explained by the PC than by the original predictors? And, if the answer to this question is yes, then 2. Are the leading PC more able to explain the response than the rest of the PC?

Question 2, as mentioned above, was posed from the beginning in order to determine the usefulness of PCA in learning settings. And this question has been answered recently by Artemiou and Li ([1]). They proved that if there are p PC, and $1 \leq i \leq j \leq p$, then the probability of the response being closer to the i-th PC than to the j-th PC is higher than the probability of the complementary event —the response being closer to the j-th PC than to the i-th PC. However, Question 2 only is important if Question 1 has positive answer, because if the original input variables explain better the answer than the PC, then a PCA is not justified. Therefore, we focus our interest here in Question 1. The aim here is to propose an explanation for the answer to this question showing, in general, that the probability of all the PC explaining better (than all the original variables) a response variable is high even in low dimensions, and increases exponentially as p grows to infinite.

The main result of this article is Theorem 2. It answers partially Question 1 on why PC usually do a better job explaining a response than the original input variables. Theorem 2 is in fact a particular version of Theorem 1, a geometric result that establishes the following: let us assume we have a given orthogonal base of a p-dimensional vectorial space, assume also that we select according to a uniform distribution a p-dimensional base for the same space (so it does not have two orthogonal vectors a.s., and it generates the space a.s.). Then, if we select a new vector uniformly at random, this new vector is going to be closer (in terms of angular distance) to the orthogonal base than to the random base with probability no less than $1/2$ for every p and this probability converges approximately to 0.6826 as $p \to \infty$.

2 Setting

Consider the space $\mathbb{R}^p$ with orthogonal base $(u_1, \ldots, u_p)$. Take Ω to be the space of all unidimensional lines traversing the origin. Select from Ω uniformly and independently a set of $p+1$ lines $v_1, \ldots, v_p$ and w. Note that $v_1, \ldots, v_p$ form almost surely a basis for $\mathbb{R}^p$. Here *uniformly* means that we have a unit closed ball centered around the origin $B[0,1]$ with surface $\tilde{B}$. Then, calling $\mathcal{L}_p$ the Lebesgue measure in p dimensions, for every $A \subset \tilde{B}$, we have that the probability of an element in Ω to traverse A is given by

$$\frac{\mathcal{L}_{p-1}(A \cup -A)}{\mathcal{L}_{p-1}\tilde{B}}, \tag{1}$$

where $-A = \{x \in \tilde{B} : -x \in A\}$.

Let $\tilde{\phi}_i$ be the angle between w and u_i, and define $\phi_i := \min\left\{\left|\tilde{\phi}_i\right|, \left|\pi - \tilde{\phi}_i\right|\right\}$. Let $\tilde{\varphi}_i$ be the angle between w and v_i, and define $\varphi_i := \min\{|\tilde{\varphi}_i|, |\pi - \tilde{\varphi}_i|\}$. Let $\tilde{\theta}_{ij}$ be the angle between w_j and v_i, and define $\theta_{ij} := \min\left\{\left|\tilde{\theta}_{ij}\right|, \left|\pi - \tilde{\theta}_{ij}\right|\right\}$.

Call $\mathcal{O}_w$ the orthant (hyper-octant) where w is located in, and, analogously, $\mathcal{O}_{u_i}$ and $\mathcal{O}_{v_j}$ the orthants where u_i and v_j live in, respectively. For convenience, let us refer to the orthants by the combination of signs they correspond to (i.e., when $p = 2$ the four quadrants can be denoted by $(+,+), (+,-), (-,+), (-,-)$).

Figure 1 gives a geometric representation of expectations of uniformly and independently chosen random variables (v_1, v_2) for $p = 2$. Figure 2 gives a geometric representation of realizations of uniformly and independently chosen random variables (v_1, v_2) and w for $p = 2$ illustrating how which set, (v_1, v_2) or (u_1, u_2), explains better the r.v. w in terms of angular distance.

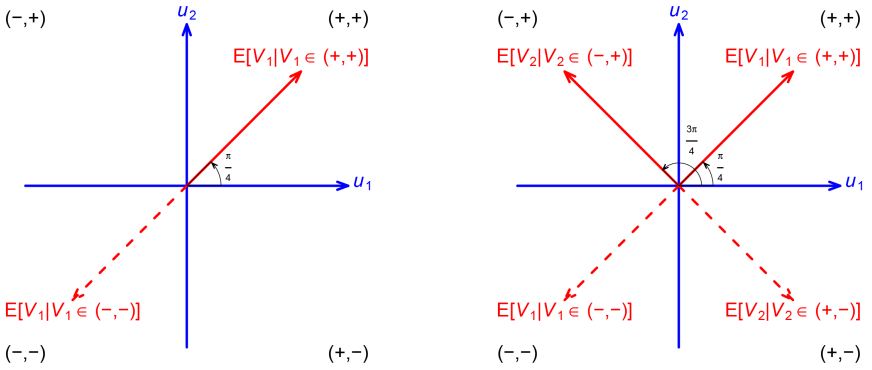

Figure 1: Geometric representation for $p = 2$ of the conditional expectation of v_1 and v_2 (red vectors) given the quadrant in which they are living. Blue vectors represent u_1 and u_2.

We prove the following result:

Theorem 1. *1. For k fixed and $i \in \{1, \ldots, p\}$, the probability of w being closer to u_k than to v_i is bigger than or equal 1/2.*

2. For k fixed and $i \in \{1, \ldots, p\}$, the probability of w being closer to u_k than to v_i converges to $\Phi(1) - \Phi(-1) \approx 0.6826$.

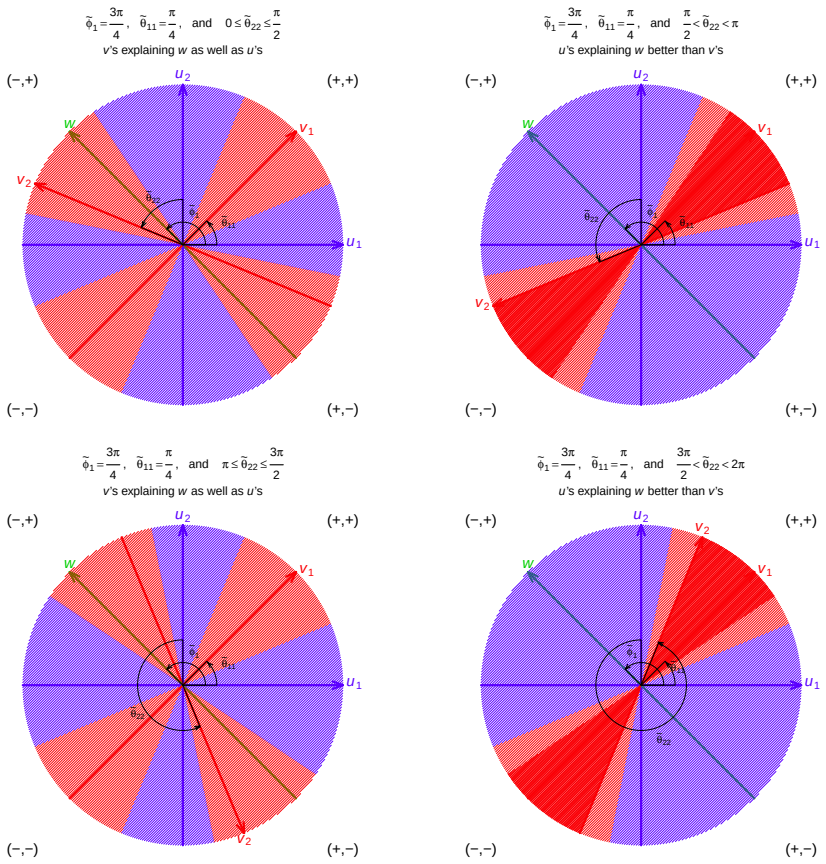

Figure 2: Geometric representation for $p = 2$ of v_1, v_2 and w. Blue vectors represent u_1 and u_2, red vectors represent realizations of v_1 and v_2, and the green vector represents w. Areas of circle sectors of u explaining w better than v are shaded in blue, and sectors of v explaining w better than u are shaded in red. The overlapping sectors of variables v appear in darker red. Notice that when "u explains w better than v", then the red shaded regions of v_1 and v_2 overlap. Also, notice that the case of "u explaining w as well as v" includes the situation where $(v_1,\ v_2)$ are mutually orthogonal. In all plots, illustrations are for an arbitrary realization of w : $\tilde{\phi}_1 = \widehat{(u_1, w)} = \frac{3\pi}{4}$, as well as v_1 : $\tilde{\theta}_{11} = \widehat{(u_1, v_1)} = \frac{\pi}{4}$, and any realizations of v_2 : $\tilde{\theta}_{22} = \widehat{(u_2, v_2)}$ in the indicated intervals.

3 Proof of Theorem 1

Call E_w the event "w is located in the orthant $\mathcal{O}_w$" and, analogously, let E_{v_i} as the event "v_i is living in the orthant $\mathcal{O}_{v_i}$". Because of (1) these events have

each probability 2^{1-p}. Assume without loss of generality that $\mathcal{O}_w$ is $(+,\cdots,+)$. Then

$$\mathbf{E}[w|E_w] = (1,\ldots,1)p^{-1/2},$$

and from this we obtain

$$\mathbf{E}[\cos\phi_k|E_w] = p^{-1/2}, \tag{2}$$

for $k = 1,\ldots,p$. Analogously to (2), we have that

$$\mathbf{E}[\cos\theta_{ik}|E_{v_i}] = p^{-1/2},$$

for $i,k = 1,\ldots,p$. Now, for $0 < j \leq p/2$, let j be the number of different signs between $\mathcal{O}_w$ and $\mathcal{O}_{v_i}$. Then we obtain that

$$\mathbf{E}\left[\cos\varphi_i|E_w,E_{v_i}\right] = 1 - 2jp^{-1}, \tag{3}$$

and, from (2) and (3),

$$\mathbf{E}\left[\cos\phi_k - \cos\varphi_i|E_w,E_{v_i}\right] = p^{-1/2} + 2jp^{-1} - 1, \tag{4}$$

so we get that (4) is nonnegative iff $\left\lceil \frac{p-p^{1/2}}{2} \right\rceil \leq j \leq \lfloor p/2 \rfloor$. Call $p^* := \left\lceil \frac{p-p^{1/2}}{2} \right\rceil$. Then, for $G_{ik}^+ := \{\cos\phi_k - \cos\varphi_i \geq 0\}$ we get that,

$$\mathbf{1}\left\{G_{ik}^+|E_w,E_{v_i}\right\} = \begin{cases} 1 & \text{if } p^* \leq j \leq \lfloor d/2 \rfloor \\ 0 & \text{if } 0 \leq j \leq p^*-1, \end{cases} \tag{5}$$

therefore

$$\mathbf{E1}\left\{G_{ik}^+|E_w,E_{v_i}\right\} = \begin{cases} 2^{1-p}\sum_{j=p^*}^{\lfloor p/2 \rfloor}\binom{p}{j} & \text{if } p \text{ is odd,} \\ 2^{1-p}\left(\sum_{j=p^*}^{p/2-1}\binom{p}{j} + \frac{\binom{p}{p/2}}{2}\right) & \text{if } p \text{ is even.} \end{cases} \tag{6}$$

As a visualization of the rate of convergence, we show in Figure 3 the corresponding convergence plot of $\mathbf{E1}\left\{G_{ik}^+|E_w,E_{v_i}\right\}$ as a function of p.

Now, by (1), total probability and independence,

$$\mathbf{P}\left[G_{ik}^+\right] = \mathbf{E1}\left\{G_{ik}^+|E_w,E_{v_i}\right\}. \tag{7}$$

Proposition 1. *$\boldsymbol{P}\left[G_{ik}^+\right] \geq 1/2$, for all $p \geq 2$.*

Some comments are in order before proceeding. Because of uniformity we are only considering half of the space (2^{p-1} orthants), since for all $\omega \in \Omega$, if ω lives in a given orthant $\mathcal{O}_\omega$, then it is also living in $-\mathcal{O}_\omega$. From this observation we see why all the combinations considered in (6) are located on the left half of the Pascal triangle for each p. Call $T(p)$ the total number of terms in this half of the Pascal triangle; then we have $T(p) = \lfloor p/2 \rfloor$. For a given p, call $R(p)$ the number of terms being added in (6), and we have that $R(p) = \lfloor p/2 \rfloor - p^* + 1$. Call also $L(p) = T(p) - R(p)$. Finally, let $D(p)$ be $\mathbf{P}\left[G_{ik}^+\right] - \mathbf{P}\left[\left(G_{ik}^+\right)^c\right]$, then the statement in Proposition 1 is equivalent to prove that $D(p) \geq 0$ and it is enough to look at the numerators of these probabilities.

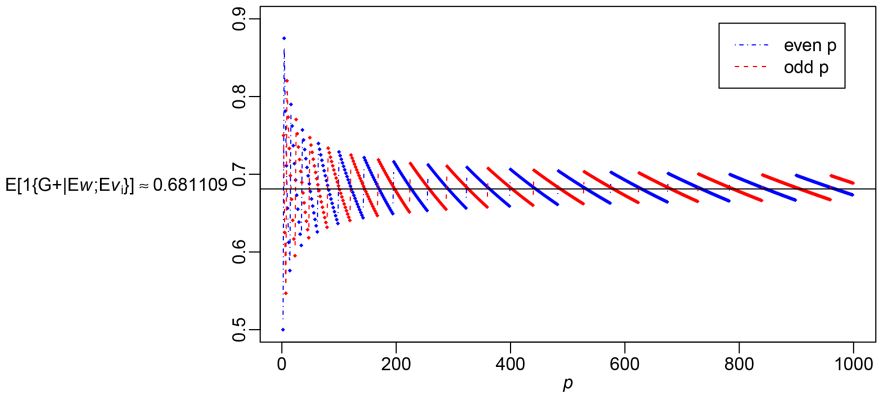

Figure 3: Convergence plot of **E1** $\left\{G_{ik}^{+}|E_w, E_{v_i}\right\}$ for $p \in \{2, \ldots, 993\}$.

Lemma 1. *$L(p)$, $C(p)$, $T(p)$ and $D(p)$ satisfy the following properties:*

(1) $T(p) = p/2$ *for all even* p.

(2) $T(p+1) = T(p)$ *for all even* p.

(3) $T(p+2) = T(p) + 1$ *for all* p.

(4) L *is increasing and either* $L(p+1) \in \{L(p), L(p)+1\}$.

(5) $R(p+1) \in \{R(p)-1, R(p), R(p)+1\}$

(6) $R(p+2) \in \{R(p), R(p)+1\}$.

(7) $L(p+2) \in \{L(p), L(p)+1\}$.

(8) Let $p \geq 2$ *and define* $\mathcal{L}_l := \{p \in \mathbb{N} : L(p) = l\}$. *Then* $\#\mathcal{L}_l$, *the cardinality of* $\mathcal{L}_l$, *is at least 2 and at most 3 for all* $p \geq 2$.

(9) For $p, p' \in \mathcal{L}_l$ *and* $p < p'$, *we have that* $2^p D(p) < 2^{p'} D(p')$.

Proof. (1), (2) and (3) follow from the definitions. (4) follows from the fact that $\{p^*\}$ is an increasing sequence and $(p+1)^* - p^* \leq 1$. The following ones can be easily deducted from the former properties. □

Note that $D(p)$ is not monotonically increasing. When $L(p+1) = L(p)+1$, we have that $R(p+1)$ is either $R(p)$ or $R(p)-1$. For such cases, $D(p+1) < D(p)$. However, the next proposition proves that despise this fact, the subsequence $\{\min_{p \in \mathcal{L}_l} D(p)\}_{l=0}^{\infty}$ is increasing.

Lemma 2. $S(\min \mathcal{L}_l) > S(\min \mathcal{L}_l - 2)$, *for all* $l > 0$.

Proof. First, note that for all $l \geq 1$, we have that $L(\min \mathcal{L}_l) = L(\min \mathcal{L}_l - 2) + 1$. Thus, $R(\min \mathcal{L}_l) = R(\min \mathcal{L}_l - 2)$ (see Properties (3), (6) and (7) in Lemma 1). Using this fact and basic combinatorics properties, it is easily seen that $D(p+2) = 4D(p) - 2\left(\binom{p}{p^*} - \binom{p}{p^*-1}\right)$, which proves the Lemma. □

Proof of Proposition 1. Lemmas 1 and 2, together with the fact that $D(2) = D(\min \mathcal{L}_0) = 1$, prove the result. □

Proof of Theorem 1. Part 1 follows directly from Proposition 1.

To prove the Part 2, let Y_p be a $\text{Bin}(p, 1/2)$. By the central limit theorem we have:

$$\mathbf{P}\left[Y_p < (p - p^{1/2})/2\right] \to \Phi(-1),$$

therefore, by symmetry of the sequence $\binom{p}{0}, \binom{p}{1}, \ldots, \binom{p}{p}$, we obtain the result. □

As a visualization of the rate of convergence, we show in Figure 3 the corresponding convergence plot of (6) as a function of p.

Remark 1. Seeing (7) from the viewpoint of the unconditioned probability of G_{ik}^+ at the left-hand side, we are looking at the event of, given i and k, having u_k closer to w than v_i. Therefore, given k, the r.v. $U_k \sim \text{Bin}(p, \mathbf{P}[G_{ik}^+])$ tells us the probability that u_k explains better w than U_k of the v's. For instance, since $\mathbf{P}[G_{ik}^+] \approx 2\mathbf{P}[(G_{ik}^+)^c]$ (indeed, it can be shown that $\mathbf{P}[G_{ik}^+] > 2\mathbf{P}[(G_{ik}^+)^c]$ for $p \geq 783$), then $\mathbf{P}[U_k = p] \approx 2^p \mathbf{P}[U_k = 0]$. That is, the probability of u_k explaining better w than all the u's is exponentially higher than the probability of all the v's explaining better w than u_k.

Remark 2. On the other hand, the right-hand side of (7), because of the condition on E_w, implies that in expectation w is closer to *all* $u_1, \ldots, u_p$ than to v_i, since being in "the middle" of an orthant means that w is at the same distance of each u_i, $i = 1, \ldots, p$. Therefore, as in the previous Remark, for large p, Theorem 1 reveals the conditionally expected proportion of v's that do not explain w as well as *all* the u's. The value is around 2/3. The downside of this conditioning is that around 1/3 of the v's are expected to explain better w than all the u's.

Remark 3. Finally, note that having compared the distance from w to v_i and to u_k as we did above, it is difficult to add to the analysis a second orthogonal vector $u_{k'}$, with $k \neq k'$, since the angle between w and $u_{k'}$ is going to be dependent on the angle betwen w and u_k.

4 Extension to random variables

One advantage of the approach used in the former section is the relatedness of the cosine similarity and the correlation coefficient of two random variables. The only difference being that the correlation coefficient is invariant to translations, while the cosine similarity is not. From the geometrical point of view, the problem is solved by restricting ourselves to consider only the vectors traversing the origin of the p-dimensional Euclidean space, precisely in the way Ω is defined. From the probabilistic point of view, however, this restriction does not affect,

since there is no loss of generality once we consider random vectors centered around the origin of the space:

Let Ω' be the space of all the continuous random variables in $\mathbb{R}^p$ with expected value 0 and such that any p-vector $\mathbf{X}$ of p of them has a covariance matrix Σ. Let $\mathbf{Y}$ be the set of principal components of $\mathbf{X}$. Then the p lines partition the space in (almost) equal regions. Consider any continuous random variable Z living in $\mathbb{R}^p$ such that $\mathbf{E}Z = 0$. Let Ω be the space of unidimensional supports of the random variables in $\mathbb{R}^p$. We have the following result:

Theorem 2. *1. For k fixed and $i \in \{1, \ldots, p\}$, we have that*

$$\boldsymbol{P}[|\rho_{Y_k Z}| \geq |\rho_{X_i Z}|] \geq 1/2.$$

2. For k fixed and $i \in \{1, \ldots, p\}$, we have that

$$\boldsymbol{P}[|\rho_{Y_k Z}| \geq |\rho_{X_i Z}|] \to \Phi(1) - \Phi(-1) \approx 0.6826.$$

Proof. Take $S_{Y_1}, \ldots, S_{Y_p}$ to be the (unidimensional) supports of $Y_1 \ldots, Y_p$, respectively. Take $S_{X_1}, \ldots, S_{X_p}$ the (unidimensional) supports of $X_1, \ldots, X_p$, respectively. Take S_Z to be the support of Z. Then the result follows from Theorem 1. □

Some comments are in order here:

- In the context of geometry, Theorem 1 is true for every set of orthogonal vectors, but their being orthogonal through uniform random selection is not happening a.s. However, even if $v_1, \ldots, v_p$ are arbitrarily selected to be orthogonal, then it can be seen that they explain the space as good as $u_1, \ldots, u_p$, but no better.

- Because of the former point, the uncorrelated random variables do not need to be the set of principal components. However, the condition of knowing the covariance matrix Σ allows us to identify a particular set of mutually orthogonal random variables for which we know explicitly the joint distribution and the marginals.

- When it comes to PCA we are not doing two things: First, we are not using the variance of the PC in any sense in this proof, except to find the directions of the components themselves. To be sure, in [1], the authors used the variance of the principal components to prove that it is more probable that the response Z is closer to the principal component Y_i than to the principal component Y_j, for $i < j$.

- Second, we are not projecting the space to any lower dimension when using principal components. As it is known, reducing the space in contexts of regression might be problematic (see for instance [4] and [5]). But Theorem 1 is true for all the p PC in the space.

- Remarks 1 and 2 also admit natural extensions to this scenario of random variables and principal components. So we get, for instance, that the probability of *all* principal components explaining better the response is exponentially higher than the probability of *all* the original input explaining it better. Of course, not even principal components are immune to the curse of dimensionality, but it shows that it affects the PC explanatory power at a lower speed.

5 Discussion

Principal components analysis (PCA) is a widely used technique but some mystery has remained as to why it's a reasonable thing to do when modeling a response-predictor relationship since it was introduced in [7] and [8] about 60 years ago. We offer a partial answer to this question. A more comprehensive answer will have to look at the probability of the response being closer to at least one of the PC than to all of the original input variables, but this is still an open problem.

Acknowledgements: All authors supported in part by NIH grant NCI R01-CA160593A1. We would like to thank Rob Tibshirani, Steve Marron and Hemant Ishwaran for helpful discussions of the work. DD would like to thank Juan Saenz for his suggestions on how to prove the results and Federico Ardila for an early very useful comment on Proposition 1.

References

[1] Artemiou A., Li B. (2009) On Principal Components and Regression: A Statistical Explanation of a Natural Phenomenom. *Statistica Sinica* **19** 1557–1565.

[2] Cook R. D. (2007) Fisher Lecture: Dimension reduction in regression. *Statist. Sci.* **22** 1–40.

[3] Cox D. R. (1968) Notes on some aspects of regression analysis. *J. Roy. Ststist. Soc. Ser. A*. **131** 265–279.

[4] Joliffe I. (1982) A Note on the Use of Principal Components in Regression. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* **31** (3), 300–303.

[5] Hadi S., Ling R.F. (1998) Some Cautionary Notes on the Use of Principal Components Regression. *The American Statistician* **52** (1), 15–19.

[6] Hocking R. R. (1976) The analysis and selection of variables in linear regression. *Biometrics* **32** 1–49.

[7] Hotelling H. (1957) The relationship of the newer multivariate statistical methods to factor analysis. *Brit. J. Stat. Psychol.* **10** 69–70.

[8] Kendall M. G. (1957) *A Course in Multivariate Analysis.* Griffin.

[9] Mosteller F., Tuckey J. W. (1977) *Data Analysis and Regression.* Addison-Wesley.

[10] Scott D. (1992) *Multivariate density estimation.* Wiley.